



Europäisches Patentamt
European Patent Office
Office européen des brevets



⑪ Publication number:

0 589 219 A2

12

EUROPEAN PATENT APPLICATION

② Application number: 93113467.0

⑤1 Int. Cl.⁵: G11B 27/00, G11B 20/10,
G06F 3/06

zz Date of filing: 24.08.93

③ Priority: 23.09.92 US 949673

(43) Date of publication of application:
30.03.94 Bulletin 94/13

⑧ Designated Contracting States:
DE FR GB

⑦ Applicant: **International Business Machines Corporation**
Old Orchard Road
Armonk, N.Y. 10504(US)

72) Inventor: Oded, Cohn
Gut Levin 20/4
Haifa 32922(IL)
Inventor: Howard, Hartung Michael
8040 East Alteza Vista
Tucson, Arizona 85715-2848(US)
Inventor: Frank, Micka William

Mendeli 6A
Haifa 32447(IL)
Inventor: **Yoram, Novick**
14 Got Levin Street No.28,
Ramot Sapir
Haifa 32922(IL)
Inventor: **Nimrod, Rotics**
14 Shunamit Street
Haifa 34562(IL)
Inventor: **Alexander, Winokur**
23 Dryfuss Street
Haifa 35434(IL)
Inventor: **Efraim, Zeidner**
Mivtza Kadesh
Kiryat-chaim 26272(IL)

74 Representative: Bravi, Alfredo et al
c/o IBM SEMEA S.p.A
Direzione Brevetti
MI SEG 024
P.O. Box 137
I-20090 Segrate (Milan) (IT)

54 Method and system for non-specific data retrieval in a data processing system.

57 A method and system for efficiently accessing desired datasets among multiple datasets which are stored at specific data addresses within multiple storage systems (20-26) which are coupled to a host system (12) via a storage subsystem controller (18). A data request is transmitted from the host system to the storage subsystems via the data channel. The data request specifies non-address attributes for desired datasets, such as boundary addresses for large data extents including many datasets or a request for all datasets modified since the occurrence of a specified event. The data request is then processed at the storage subsystem controller (18) to determine a data address for each dataset within the storage system (20-26) which possess the desired attributes. Thereafter, the desired datasets are transmitted to the host system (12) in association with a specific address for each dataset. A selected status

message is transmitted from the storage subsystem controller when no more datasets are located which possess the desired attributes.

The present invention relates in general to improved methods and systems for managing datasets stored within storage subsystems in a data processing system and in particular to improved methods and systems for accessing desired datasets stored within storage subsystems in a data processing system. Still more particularly, the present invention relates to improved methods and systems for accessing desired datasets within a storage subsystem utilizing non-address attributes set forth within a data request.

Data processing systems frequently include large scale storage devices, such as Direct Access Storage Devices (DASD) which are located externally to a host computer system and sometimes at significant distances from the host computer system. Communication from the host computer system to the DASD is typically accomplished over signal cables, called "data channels", extending between the DASD and its control unit and connecting the DASD devices to the host computer system.

Current technology provides DASD units with several separate disks, all rotating on a single spindle. These disks or platters are accessed by head disk assemblies with a transducing head providing access to one surface of each disk. There may be, for example, nine platters in a disk drive providing sixteen usable surfaces with one of the usable surfaces used for maintaining accurate tracking capability. In such units there are fifteen usable surfaces for data and when all heads are correctly positioned a cylinder of fifteen physical, recording tracks may be accessed.

DASD units frequently use a so-called "Count-Key-Data" architecture (CKD) where records written on a track within a DASD unit are provided with a count field (an identification), a key length field and a data field. A record may occupy one or more units of real storage. A "dataset" is a logical collection of multiple records which may be stored on contiguous units of real storage or which may be dispersed. Data is then stored and/or retrieved from a DASD using write and read requests which are issued by the host system. The mechanism which enables host systems to retrieve data which has previously been stored on a disk is the "data disk address". Therefore, when issuing a write request, the host system specifies where on the DASD storage subsystem the data should be placed. Later, if the host system wishes to retrieve this data it will issue a read request utilizing the same address.

Thus, data stored on a disk within a storage subsystem is always associated with a unique data descriptor which identifies that data. In a write request the host specifies the data descriptor together with the data to be stored. In a read request

the host specifies the data descriptor of the data it wishes to receive. In response to such a request the DASD subsystem will send the referenced data back to the host. For purposes of explanation herein such read requests which utilize data descriptors are referred to as "specific read requests".

Those skilled in the art will appreciate that it would be advantageous to permit a host system to retrieve data from a DASD subsystem on a basis other than the data descriptor. For example, in cases where the host system requests a large amount of data occupying many disk tracks the efficiency of transferring that data to the host might be enhanced if the order of the transfer is adjusted in order to minimize both seek time and latency time within the data storage subsystem. This is not generally possible since host systems do not know the head location within the DASD subsystem and thus are not able to issue specific read requests which would minimize the disk seek and latency time.

Additionally, many systems exist which are utilized to create so-called "backup" copies of data. In an incremental backup copy only that data which has been modified since the previous copy need be transferred to the host system. In such cases, the host does not initially know what data has been modified and thus the host can not issue specific read requests without issuing a query to the storage subsystem to determine the data descriptors for the data which has been updated since the previous copy. In view of the above, those skilled in the art will appreciate that it would be desirable to permit a host system to retrieve data from a storage subsystem by specifying certain attributes of the data rather than the actual data address.

It is therefore one object of the present invention to provide an improved method and system for managing datasets stored within storage subsystems in a data processing system.

It is another object of the present invention to provide an improved method and system for accessing desired datasets stored within storage subsystems in a data processing system.

It is yet another object of the present invention to provide an improved method and system for accessing desired datasets within a storage subsystem utilizing non-address attributes set forth within a data request.

The foregoing objects are achieved by the invention as claimed. The method and system of the present invention may be utilized to efficiently access desired datasets among multiple datasets which are stored at specific data addresses within multiple storage subsystems coupled to a host system via a storage subsystem controller and a data channel. A data request is transmitted from the host system to the storage subsystems via the

data channel. The data request specifies non-address attributes for desired datasets, such as boundary addresses for large data extents including many datasets or a request for all datasets modified since the occurrence of a specified event. The data request is then processed at the storage subsystem controller to determine a data address for each dataset or portions thereof within the storage subsystem which possess the desired attributes. Thereafter, the desired datasets or portions thereof are transmitted via the data channel to the host system in association with a specific address for each dataset or portion thereof. A selected status message is transmitted from the storage subsystem controller when no more datasets are located which possess the desired attributes. In this manner, the retrieval of data from a storage subsystem is greatly enhanced.

The above as well as additional objects, features, and advantages of the present invention will become apparent in the following detailed written description.

The novel features believed characteristic of the invention are set forth in the appended claims. The invention itself however, as well as a preferred mode of use, further objects and advantages thereof, will best be understood by reference to the following detailed description of an illustrative embodiment when read in conjunction with the accompanying drawings, wherein:

Figure 1 is a pictorial representation of a data processing system which may be utilized to access desired datasets in accordance with the method and system of the present invention;

Figure 2 is a high level logic flowchart illustrating a non-specific data request in accordance with the method and system of the present invention;

Figure 3 is a high level logic flowchart illustrating a response to a non-specific data request in accordance with the method and system of the present invention.

With reference now to the figures and in particular with reference to Figure 1, there is illustrated a pictorial representation of a data processing system 10 which may be utilized to access desired datasets in accordance with the method and system of the present invention. As illustrated, data processing system 10 includes a host computer system 12, which is coupled to a storage system which comprises a plurality of Direct Access Storage Devices (DASD) 20, 22, 24, and 26 via a data channel 16 and a storage subsystem controller 18. Applications 14 within host computer system 12 may be utilized to access and manipulate data stored within the storage subsystem in a manner well known in the art.

Access to desired datasets within the storage subsystem which is comprised of storage subsystem controller 18, DASD 20, DASD 22, DASD 24, and DASD 26 is typically accomplished in the prior art, by setting forth and specifying a unique address or data descriptor for the desired data and transmitting that data descriptor to storage subsystem controller 18 via data channel 16. Data channel 16 may constitute an electrical cable or, in a modern state-of-the-art data processing system may be implemented utilizing a fiber optic cable, such as the International Business Machines Enterprise System Connection (ESCON). Data channel 16 is coupled to the upper port of storage subsystem controller 18 in a manner well known in the art.

Those skilled in the art will appreciate that storage subsystem controller 18 may be implemented utilizing any state-of-the-art storage subsystem controller such as the International Business Machines Corporation Model 3990. Storage subsystem controller 18, in the depicted embodiment of the present invention, includes a memory 28 which may be utilized to store data within storage subsystem controller 18 prior to transmittal of that data to host computer system 12 via data channel 16. Thus, in accordance with the prior art, a specific request by host computer system 12 to read data within the storage subsystem is accomplished by transmitting a specific data address for that data to storage subsystem controller 18, via data channel 16 and retrieving that data from a selected DASD via the lower port of storage subsystem controller 18. That data is then transmitted to host computer system 12 in response to the read request.

As described above, it should be apparent upon reference to the foregoing that the efficiency of retrieval of data from the multiple Direct Access Storage Devices within the storage subsystem may be greatly enhanced by permitting host computer system 12 to retrieve data therefrom utilizing a non-specific read request. That is, a request for data which specifies that data utilizing a non-address attribute. For example, in the transfer of large amounts of data, wherein the order in which the data has been transferred does not matter, the efficiency of the transfer may be greatly enhanced by minimizing seek and latency delays by permitting the storage subsystem to retrieve that data which is physically closest to the head within the Direct Access Storage Device.

This may be accomplished in accordance with the method and system of the present invention by specifying boundary addresses between which all data present is to be retrieved by the storage subsystem controller 18 and transmitted to the host computer system 12. Thus, all data records be-

tween those boundary addresses may be retrieved, without specifying exact address data for each desired dataset. As utilized within the present application, the term "non-address attribute" may include the specification of boundary addresses and still be considered a "non-address attribute" since that attribute does not include specific addresses for datasets contained therein.

Similarly, the retrieval of all data which has been modified since a previous backup copy was created may be retrieved by specifying as the non-address attribute those datasets which have been updated subsequent to a specified event. These two examples are merely illustrative of the process by which specific datasets may be retrieved by specifying non-address attributes in the manner described herein, so long as the attribute is one which is recognized by host computer system 12 and storage subsystem controller 18. Of course, non-address attributes may be transmitted from host computer system 12 or may be predetermined and stored within storage subsystem controller 18 or within a DASD in association with the datasets themselves.

Referring now to Figure 2, there is depicted a high level logic flowchart which illustrates the generation of a non-specific data request in accordance with the method and system of the present invention. As depicted, the process begins at block 40 and thereafter passes to block 42. Block 42 illustrates a determination of the non-address attribute for the desired datasets. As described above, this non-address attribute may constitute boundary addresses between which all data is to be transferred to the host computer system, or the attribute may constitute all datasets which have been updated subsequent to a specified event within data processing system 10. Thereafter, the process passes to block 44.

Block 44 illustrates the formulation of a non-specific data request at the host. In the depicted embodiment of the present invention this is accomplished by transmitting a specific command to the storage subsystem controller which is previously identified within data processing system 10 as requesting transfer of all data which possesses an attribute which is specified within a field associated with that command. Next, the non-specific data request is transmitted to the storage subsystem controller for processing at the storage subsystem controller. The process of generating a non-specific data request then terminates, as illustrated at block 48.

Referring now to Figure 3, there is depicted a high level logic flowchart which illustrates a response to a non-specific data request at the storage subsystem controller in accordance with the method and system of the present invention. As

depicted, the process begins at block 60 and thereafter passes to block 62. Block 62 illustrates a determination of whether or not a non-specific data request has been received from the host computer system. If not, the process merely iterates until such time as a non-specific data request has been received.

In the event a non-specific data request of the format described above is received from the host computer system, as depicted at block 62, the process passes to block 64. Block 64 illustrates the location by the storage subsystem controller of a dataset which possesses the specified non-address attributes. The process then passes to block 66 which depicts a determination of whether or not a dataset possessing the specified non-address attributes has been located and if not, the process passes to block 68. Block 68 illustrates the transmittal of a status message to the host computer system which indicates "NO DATASETS FOUND" and the process then passes to block 70 and returns.

In this manner, those skilled in the art will appreciate upon reference to the foregoing that by transmitting a non-specific read request from the host computer system to the storage subsystem controller and by permitting the storage subsystem controller to locate datasets within the storage subsystems which possess a specified non-address attribute the efficiency of transferring data to the host computer system may be greatly enhanced, not only by obviating the necessity of transmitting specific data descriptors from the host computer system but also by making it possible to retrieve data where the specific data descriptors are unknown or not determinable.

Referring again to block 66, in the event a dataset possessing the specified non-address attributes has been located, block 72 illustrates the transmittal of that dataset and the dataset address for that dataset to the host computer system. This is an important feature of the present invention since the transmittal of that dataset to the host computer system without including the dataset address will not permit the host computer system to locate that dataset at a subsequent time. Next, the process passes to block 74 which illustrates a determination of whether or not any more datasets have been located which possess the specified non-address attributes contained within the non-specific data requests. If so, the process returns iteratively to block 72 and once again transmits a located dataset and that dataset's address to the host computer system.

Referring again to block 74, in the event no more datasets have been located which possess the specified non-address attributes contained within the non-specific read request, the process

passes to block 76. In a manner similar to that described above with respect to block 68, block 76 illustrates the transmitting of a status message to the host computer system which states "NO MORE DATASETS FOUND" and the process then returns, as depicted at block 70.

Upon reference to the foregoing those skilled in the art will appreciate that the Applicants herein have created a data processing system which permits the highly efficient recovery of desired datasets from among a plurality of datasets stored within a storage subsystem without requiring the host computer system to specify specific address locations for those desired datasets. By permitting the host computer system to transmit a non-specific data request which includes a specification of desired non-address attributes for datasets within the storage subsystem, the host computer system may rapidly and efficiently retrieve data stored within the storage subsystem in the manner described herein.

While the invention has been particularly shown and described with reference to a preferred embodiment, it will be understood by those skilled in the art that various changes in form and detail may be made therein without departing from the spirit and scope of the invention.

Claims

1. A method in a data processing system (10) for efficiently accessing desired datasets among a plurality of datasets which are stored at specific data addresses within a plurality of storage systems (20-26) coupled to a host system (12) via a storage subsystem controller (18) and a data channel (16), said method comprising the steps of:

transmitting a data request (42-46) from said host system (12) to said plurality of storage systems (20-26) via said data channel (16), said data request specifying a non-address attribute of said desired datasets;

processing said data request (62-66) at said storage subsystem controller (18) to determine a data address for each dataset among said plurality of datasets which possesses said non-address attribute; and

transmitting each dataset (72) possessing said non-address attribute to said host system (12) via said data channel (16).

2. The method in a data processing system for efficiently accessing desired datasets according to Claim 1, wherein said step of transmitting each dataset (72) possessing said non-address attribute to said host system (12) via said data channel (16) further comprises the

step of transmitting a specific data address in association with each dataset possessing said non-address attribute.

- 5 3. The method in a data processing system for efficiently accessing desired datasets according to Claim 1 or 2, wherein said step of transmitting a data request (42-46) from said host system (12) to said plurality of storage systems (20-26) via said data channel (16) comprises the step of transmitting a data request for all datasets modified subsequent to a specified event.
- 10 4. The method in a data processing system for efficiently accessing desired datasets according to Claim 1 or 2, wherein said step of transmitting a data request (42-46) from said host system (12) to said plurality of storage systems (20-26) via said data channel (16) comprises the step of transmitting a data request for all datasets which have been stored between a first data address and a second data address.
- 15 5. The method in a data processing system for efficiently accessing desired datasets according to any preceding Claims, further including the step of transmitting a selected status message from said storage subsystem controller (18) to said host system (12) in response to a failure of said storage subsystem controller (18) to locate additional datasets possessing said non-address attribute.
- 20 30 35 40 45 50 55 60 65 70 75 80 85 90
- 6. A data processing system (10) for efficiently accessing desired datasets among a plurality of datasets which are stored at specific data addresses within a plurality of storage systems (20-26) coupled to a host system (12) via a storage subsystem controller (18) and a data channel (16), said data processing system comprising:
 - means for transmitting a data request (42-46) from said host system (12) to said plurality of storage systems (20-26) via said data channel (16), said data request specifying a non-address attribute of said desired datasets;
 - means for processing (62-66) said data request at said storage subsystem controller (18) to determine a data address for each dataset among said plurality of datasets which possesses said non-address attribute; and
 - means for transmitting each dataset (72) possessing said non-address attribute to said host system (12) via said data channel (16).

7. The data processing system for efficiently accessing desired datasets according to Claim 6, wherein said means for transmitting each dataset (72) possessing said non-address attribute to said host system (12) via said attribute to said host system (12) via said data channel (16) further comprises means for transmitting a specific data address in association with each dataset possessing said non-address attribute.

5

8. The method in a data processing system for efficiently accessing desired datasets according to Claim 6 or 7, wherein said means for transmitting (42-46) a data request from said host system (12) to said plurality of storage systems (20-26) via said data channel (16) comprises means for transmitting a data request for all datasets which have been modified subsequent to a specified event.

10

15

9. The method in a data processing system for efficiently accessing desired datasets according to Claim 6 or 7, wherein said means for transmitting (42-46) a data request from said host system (12) to said plurality of storage systems (20-26) via said data channel (16) comprises means for transmitting a data request for all datasets stored between a first data address and a second data address.

20

25

10. The method in a data processing system for efficiently accessing desired datasets according to Claim from 6 to 9, further including means for transmitting a selected status message (76) from said storage subsystem controller (18) to said host system (12) in response to a failure of said storage subsystem controller (18) to locate additional datasets possessing said non-address attribute.

30

35

40

45

50

55




Fig. 1




Fig. 2





Fig. 3



(19)

Europäisches Patentamt

European Patent Office

Office européen des brevets



(11)

EP 0 589 219 A3

(12)

EUROPEAN PATENT APPLICATION

(88) Date of publication A3:
10.01.1996 Bulletin 1996/02

(51) Int. Cl.⁶: G11B 27/00, G11B 20/10,
G06F 3/06, G06F 11/14,
G06F 17/30

(43) Date of publication A2:
30.03.1994 Bulletin 1994/13

(21) Application number: 93113467.0

(22) Date of filing: 24.08.1993

(84) Designated Contracting States:
DE FR GB

• Frank, Micka William
Tucson, Arizona 85718 (US)

(30) Priority: 23.09.1992 US 949673

• Michael, Nagin Kenneth

(71) Applicant: International Business Machines
Corporation
Armonk, N.Y. 10504 (US)

Haifa 32447 (IL)

(72) Inventors:
• Oded, Cohn
Haifa 32922 (IL)
• Howard, Hartung Michael
Tucson, Arizona 85715-2848 (US)

• Yoram, Novick

Haifa 32922 (IL)

• Nimrod, Rotics

Haifa 34562 (IL)

• Alexander, Winokur

Haifa 35434 (IL)

• Efraim, Zeidner

Kiryat-chaim 26272 (IL)

(74) Representative: Lettieri, Fabrizio et al
I-20059 Vimercate (MI) (IT)

(54) Method and system for non-specific data retrieval in a data processing system

(57) A method and system for efficiently accessing desired datasets among multiple datasets which are stored at specific data addresses within multiple storage systems (20-26) which are coupled to a host system (12) via a storage subsystem controller (18). A data request is transmitted from the host system to the storage subsystems via the data channel. The data request specifies non-address attributes for desired datasets, such as boundary addresses for large data extents including many datasets or a request for all datasets modified since the occurrence of a specified event. The data request is then processed at the storage subsystem controller (18) to determine a data address for each dataset within the storage system (20-26) which possess the desired attributes. Thereafter, the desired datasets are transmitted to the host system (12) in association with a specific address for each dataset. A selected status message is transmitted from the storage subsystem controller when no more datasets are located which possess the desired attributes.

EP 0 589 219 A3



European Patent
Office

EUROPEAN SEARCH REPORT

Application Number

EP 93 11 3467

DOCUMENTS CONSIDERED TO BE RELEVANT			
Category	Citation of document with indication, where appropriate, of relevant passages	Relevant to claim	CLASSIFICATION OF THE APPLICATION (Int.Cl.5)
A	EP-A-0 381 651 (IBM) 8 August 1990 ---		G11B27/00 G11B20/10 G06F3/06 G06F11/14 G06F17/30
A	EP-A-0 150 705 (IBM) 7 August 1985 -----		
TECHNICAL FIELDS SEARCHED (Int.Cl.5)			
G11B G06F			
<p>The present search report has been drawn up for all claims</p>			
Place of search	Date of completion of the search	Examiner	
THE HAGUE	26 October 1995	Benfield, A	
CATEGORY OF CITED DOCUMENTS			
X : particularly relevant if taken alone Y : particularly relevant if combined with another document of the same category A : technological background O : non-written disclosure P : intermediate document			
T : theory or principle underlying the invention E : earlier patent document, but published on, or after the filing date D : document cited in the application L : document cited for other reasons & : member of the same patent family, corresponding document			